arXiv:1411.1279v1 [cs.SI] 3 Nov 2014

Streaming, Memory Limited Algorithms for Community Detection

Se-Young Yun
MSR-Inria
23 Avenue d'Italie, Paris 75013
seyoung.yun@inria.fr

Marc Lelarge *
Inria & ENS
23 Avenue d'Italie, Paris 75013
marc.lelarge@ens.fr

Alexandre Proutiere †
KTH, EE School / ACL
Osquldasv. 10, Stockholm 100-44, Sweden
alepro@kth.se

Abstract

In this paper, we consider sparse networks consisting of a finite number of non-overlapping communities, i.e. disjoint clusters, so that there is higher density within clusters than across clusters. Both the intra- and inter-cluster edge densities vanish when the size of the graph grows large, making the cluster reconstruction problem nosier and hence difficult to solve. We are interested in scenarios where the network size is very large, so that the adjacency matrix of the graph is hard to manipulate and store. The data stream model in which columns of the adjacency matrix are revealed sequentially constitutes a natural framework in this setting. For this model, we develop two novel clustering algorithms that extract the clusters asymptotically accurately. The first algorithm is *offline*, as it needs to store and keep the assignments of nodes to clusters, and requires a memory that scales linearly with the network size. The second algorithm is *online*, as it may classify a node when the corresponding column is revealed and then discard this information. This algorithm requires a memory growing sub-linearly with the network size. To construct these efficient streaming memory-limited clustering algorithms, we first address the problem of clustering with partial information, where only a small proportion of the columns of the adjacency matrix is observed and develop, for this setting, a new spectral algorithm which is of independent interest.

1 Introduction

Extracting clusters or communities in networks have numerous applications and constitutes a fundamental task in many disciplines, including social science, biology, and physics. Most methods for clustering networks assume that pairwise "interactions" between nodes can be observed, and that from these observations, one can construct a graph which is then partitioned into clusters. The resulting graph partitioning problem can be typically solved using spectral methods [1, 3, 5, 6, 13], compressed sensing and matrix completion ideas [2, 4], or other techniques [11].

A popular model and benchmark to assess the performance of clustering algorithms is the Stochastic Block Model (SBM) [10], also referred to as the planted partition model. In the SBM, it is assumed

*Work performed as part of MSR-INRIA joint research centre. M.L. acknowledges the support of the French Agence Nationale de la Recherche (ANR) under reference ANR-11-JS02-005-01 (GAP project).

†A. Proutiere's research is supported by the ERC FSA grant, and the SSF ICT-Psi project.

that the graph to partition has been generated randomly, by placing an edge between two nodes with probability p if the nodes belong to the same cluster, and with probability q otherwise, with $q < p$. The parameters p and q typically depends on the network size n, and they are often assumed to tend to 0 as n grows large, making the graph sparse. This model has attracted a lot of attention recently. We know for example that there is a phase transition threshold for the value of $\frac{(p-q)^2}{p+q}$. If we are below the threshold, no algorithm can perform better than the algorithm randomly assigning nodes to clusters [7, 15], and if we are above the threshold, it becomes indeed possible to beat the naive random assignment algorithm [12]. A necessary and sufficient condition on p and q for the existence of clustering algorithms that are asymptotically accurate (meaning that the proportion of misclassified nodes tends to 0 as n grows large) has also been identified [17]. We finally know that spectral algorithms can reconstruct the clusters asymptotically accurately as soon as this is at all possible, i.e., they are in a sense optimal.

We focus here on scenarios where the network size can be extremely large (online social and biological networks can, already today, easily exceed several hundreds of millions of nodes), so that the adjacency matrix A of the corresponding graph can become difficult to manipulate and store. We revisit network clustering problems under memory constraints. Memory limited algorithms are relevant in the streaming data model, where observations (i.e. parts of the adjacency matrix) are collected sequentially. We assume here that the columns of the adjacency matrix A are revealed one by one to the algorithm. An arriving column may be stored, but the algorithm cannot request it later on if it was not stored. The objective of this paper is to determine how the memory constraints and the data streaming model affect the fundamental performance limits of clustering algorithms, and how the latter should be modified to accommodate these restrictions. Again to address these questions, we use the stochastic block model as a performance benchmark. Surprisingly, we establish that when there exists an algorithm with unlimited memory that asymptotically reconstruct the clusters accurately, then we can devise an asymptotically accurate algorithm that requires a memory scaling linearly in the network size n, except if the graph is extremely sparse. This claim is proved for the SBM with parameters $p = a\frac{f(n)}{n}$ and $q = b\frac{f(n)}{n}$, with constants $a > b$, under the assumption that $\log n \ll f(n)$. For this model, unconstrained algorithms can accurately recover the clusters as soon as $f(n) = \omega(1)$ [17], so that the gap between memory-limited and unconstrained algorithms is rather narrow. We further prove that the proposed algorithm reconstruct the clusters accurately before collecting all the columns of the matrix A, i.e., it uses *less* than one pass on the data. We also propose an online streaming algorithm with sublinear memory requirement. This algorithm output the partition of the graph in an online fashion after a group of columns arrives. Specifically, if $f(n) = n^\alpha$ with $0 < \alpha < 1$, our algorithm requires as little as n^β memory with $\beta > \max\left(1-\alpha, \frac{2}{3}\right)$. To the best of our knowledge, our algorithm is the first sublinear streaming algorithm for community detection. Although streaming algorithms for clustering data streams have been analyzed [9], the focus in this theoretical computer science literature is on worst case graphs and on approximation performance which is quite different from ours.

To construct efficient streaming memory-limited clustering algorithms, we first address the problem of clustering with *partial information*. More precisely, we assume that a proportion γ (that may depend on n) of the columns of A is available, and we wish to classify the nodes corresponding to these columns, i.e., the observed nodes. We show that a necessary and sufficient condition for the existence of asymptotically accurate algorithms is $\sqrt{\gamma} f(n) = \omega(1)$. We also show that to classify the observed nodes efficiently, a clustering algorithm must exploit the information provided by the edges between observed and unobserved nodes. We propose such an algorithm, which in turn, constitutes a critical building block in the design of memory-limited clustering schemes.

To our knowledge, this paper is the first to address the problem of community detection in the streaming model, and with memory constraints. Note that PCA has been recently investigated in the streaming model and with limited memory [14]. Our model is different, and to obtain efficient clustering algorithms, we need to exploit its structure.

2 Models and Problem Formulation

We consider a network consisting of a set V of n nodes. V admits a hidden partition of K non-overlapping subsets $V_1, \ldots, V_K$, i.e., $V = \bigcup_{k=1}^K V_k$. The size of community or cluster V_k is $\alpha_k n$ for some $\alpha_k > 0$. Without loss of generality, let $\alpha_1 \leq \alpha_2 \leq \cdots \leq \alpha_K$. We assume that when the

network size n grows large, the number of communities K and their relative sizes are kept fixed. To recover the hidden partition, we have access to a $n \times n$ symmetric random binary matrix A whose entries are independent and satisfy: for all $v, w \in V$, $\mathbb{P}[A_{vw} = 1] = p$ if v and w are in the same cluster, and $\mathbb{P}[A_{vw} = 1] = q$ otherwise, with $q < p$. This corresponds to the celebrated Stochastic Block Model (SBM). If $A_{vw} = 1$, we say that nodes v and w are connected, or that there is an edge between v and w. p and q typically depend on the network size n. To simplify the presentation, we assume that there exists a function $f(n)$, and two constants $a > b$ such that $p = a\frac{f(n)}{n}$ and $q = b\frac{f(n)}{n}$. This assumption on the specific scaling of p and q is not crucial, and most of the results derived in this paper hold for more general p and q (as it can be seen in the proofs). For an algorithm π, we denote by $\varepsilon^\pi(n)$ the proportion of nodes that are misclassified by this algorithm. We say that π is asymptotically accurate if $\lim_{n\to\infty} \mathbb{E}[\varepsilon^\pi(n)] = 0$. Note that in our setting, if $f(n) = O(1)$, there is a non-vanishing fraction of isolated nodes for which no algorithm will perform better than a random guess. In particular, no algorithm can be asymptotically accurate. Hence, we assume that $f(n) = \omega(1)$, which constitutes a necessary condition for the graph to be asymptotically connected, i.e., the largest connected component to have size $n - o(n)$.

In this paper, we address the problem of reconstructing the clusters from specific observed entries of A, and under some constraints related to the memory available to process the data and on the way observations are revealed and stored. More precisely, we consider the two following problems.

Problem 1. Clustering with partial information. We first investigate the problem of detecting communities under the assumption that the matrix A is partially observable. More precisely, we assume that a proportion γ (that typically depend on the network size n) of the columns of A are known. The γn observed columns are selected uniformly at random among all columns of A. Given these observations, we wish to determine the set of parameters γ and $f(n)$ such that there exists an asymptotically accurate clustering algorithm.

Problem 2. Clustering in the streaming model and under memory constraints. We are interested here in scenarios where the matrix A cannot be stored entirely, and restrict our attention to algorithms that require memory less than M bits. Ideally, we would like to devise an asymptotically accurate clustering algorithm that requires a memory M scaling linearly or sub-linearly with the network size n. In the streaming model, we assume that at each time $t = 1, \dots, n$, we observe a column A_v of A uniformly distributed over the set of columns that have not been observed before t. The column A_v may be stored at time t, but we cannot request it later on if it has not been explicitly stored. The problem is to design a clustering algorithm π such that in the streaming model, π is asymptotically accurate, and requires less than M bits of memory. We distinguish offline clustering algorithms that must store the mapping between all nodes and their clusters (here M has to scale linearly with n), and online algorithms that may classify the nodes when the corresponding columns are observed, and then discard this information (here M could scale sub-linearly with n).

3 Clustering with Partial Information

In this section, we solve Problem 1. In what follows, we assume that $\gamma n = \omega(1)$, which simply means that the number of observed columns of A grows large when n tends to ∞. However we are typically interested in scenarios where the proportion of observed columns γ tends to 0 as the network size grows large. Let $(A_v, v \in V^{(g)})$ denote the observed columns of A. $V^{(g)}$ is referred to as the set of *green* nodes and we denote by $n^{(g)} = \gamma n$ the number of green nodes. $V^{(r)} = V \setminus V^{(g)}$ is referred to as the set of *red* nodes. Note that we have no information about the connections among the red nodes. For any $k = 1, \dots, K$, let $V_k^{(g)} = V^{(g)} \cap V_k$, and $V_k^{(r)} = V^{(r)} \cap V_k$. We say that a clustering algorithm π classifies the green nodes asymptotically accurately, if the proportion of misclassified green nodes, denoted by $\varepsilon^\pi(n^{(g)})$, tends to 0 as the network size n grows large.

3.1 Necessary Conditions for Accurate Detection

We first derive necessary conditions for the existence of asymptotically accurate clustering algorithms. As it is usual in this setting, the hardest model to estimate (from a statistical point of view) corresponds to the case of two clusters of equal sizes (see Remark 3 below). Hence, we state our information theoretic lower bounds, Theorems 1 and 2, for the special case where $K = 2$, and

$\alpha_1 = \alpha_2$. Theorem 1 states that if the proportion of observed columns γ is such that $\sqrt{\gamma}f(n)$ tends to 0 as n grows large, then no clustering algorithm can perform better than the naive algorithm that assigns nodes to clusters randomly.

Theorem 1 *Assume that $\sqrt{\gamma}f(n) = o(1)$. Then under any clustering algorithm π, the expected proportion of misclassified green nodes tends to 1/2 as n grows large, i.e., $\lim_{n\to\infty} \mathbb{E}[\varepsilon^{\pi}(n^{(g)})] = 1/2$.*

Theorem 2 (i) shows that this condition is tight in the sense that as soon as there exists a clustering algorithm that classifies the green nodes asymptotically accurately, then we need to have $\sqrt{\gamma}f(n) = \omega(1)$. Although we do not observe the connections among red nodes, we might ask to classify these nodes through their connection patterns with green nodes. Theorem 2 (ii) shows that this is possible only if $\gamma f(n)$ tends to infinity as n grows large.

Theorem 2 *(i) If there exists a clustering algorithm that classifies the green nodes asymptotically accurately, then we have: $\sqrt{\gamma}f(n) = \omega(1)$.*
(ii) If there exists an asymptotically accurate clustering algorithm (i.e., classifying all nodes asymptotically accurately), then we have: $\gamma f(n) = \omega(1)$.

Remark 3 *Theorems 1 and 2 might appear restrictive as they only deal with the case of two clusters of equal sizes. This is not the case as we will provide in the next section an algorithm achieving the bounds of Theorem 2 (i) and (ii) for the general case (with a finite number K of clusters of possibly different sizes). In other words, Theorems 1 and 2 translates directly in minimax lower bounds thanks to the results we obtain in Section 3.2.*

Note that as soon as $\gamma f(n) = \omega(1)$ (i.e. the mean degree in the observed graph tends to infinity), then standard spectral method applied on the squared matrix $A^{(g)} = (A_{vw}, v, w \in V^{(g)})$ will allow us to classify asymptotically accurately the green nodes, i.e., taking into account only the graph induced by the green vertices is sufficient. However if $\gamma f(n) = o(1)$ then no algorithm based on the induced graph only will be able to classify the green nodes. Theorem 2 shows that in the range of parameters $1/f(n)^2 \ll \gamma \ll 1/f(n)$, it is impossible to cluster asymptotically accurately the red nodes but the question of clustering the green nodes is left open.

3.2 Algorithms

In this section, we deal with the general case and assume that the number K of clusters (of possibly different sizes) is known. There are two questions of interest: clustering green and red nodes. It seems intuitive that red nodes can be classified only if we are able to first classify green nodes. Indeed as we will see below, once the green nodes have been classified, an easy greedy rule is optimal for the red nodes.

Classifying green nodes. Our algorithm to classify green nodes rely on spectral methods. Note that as suggested above, in the regime $1/f(n)^2 \ll \gamma \ll 1/f(n)$, any efficient algorithm needs to exploit the observed connections between green and red nodes. We construct such an algorithm below. We should stress that our algorithm does not require to know or estimate γ or $f(n)$.

When from the observations, a red node $w \in V^{(r)}$ is connected to at most a single green node, i.e., if $\sum_{v\in V^{(g)}} A_{vw} \leq 1$, this red node is useless in the classification of green nodes. On the contrary, when a red node is connected to two green nodes, say v_1 and v_2 ($A_{v_1 w} = 1 = A_{v_2 w}$), we may infer that the green nodes v_1 and v_2 are likely to be in the same cluster. In this case, we say that there is an *indirect edge* between v_1 and v_2.

To classify the green nodes, we will use the matrix $A^{(g)} = (A_{vw})_{v,w\in V^{(g)}}$, as well as the graph of indirect edges. However this graph is statistically different from the graphs arising in the classical stochastic block model. Indeed, when a red node is connected to three or more green nodes, then the presence of indirect edges between these green nodes are not statistically independent. To circumvent this difficulty, we only consider indirect edges created through red nodes connected to exactly two green nodes. Let $V^{(i)} = \{v : v \in V^{(r)} \text{ and } \sum_{w\in V^{(g)}} A_{wv} = 2\}$. We denote by A' the $(n^{(g)} \times n^{(g)})$ matrix reporting the number of such indirect edges between pairs of green nodes: for all $v, w \in V^{(g)}$, $A'_{vw} = \sum_{z\in V^{(i)}} A_{vz}A_{wz}$.

Algorithm 1 Spectral method with indirect edges

Input: $A \in \{0,1\}^{|V|\times|V^{(g)}|}$, V, $V^{(g)}$, K
$V^{(r)} \leftarrow V \setminus V^{(g)}$
$V^{(i)} \leftarrow \{v : v \in V^{(r)} \text{ and } \sum_{w\in V^{(g)}} A_{wv} = 2\}$
$A^{(g)} \leftarrow (A_{vw})_{v,w\in V^{(g)}}$ and $A' \leftarrow (A'_{vw} = \sum_{z\in V^{(i)}} A_{vz}A_{wz})_{v,w\in V^{(g)}}$
$\hat{p}^{(g)} \leftarrow \frac{\sum_{v,w\in V^{(g)}} A^{(g)}_{vw}}{|V^{(g)}|^2}$ and $\hat{p}' \leftarrow \frac{\sum_{v,w\in V^{(g)}} A'_{vw}}{|V^{(g)}|^2}$
$Q^{(g)}, \sigma_K^{(g)}, \Gamma^{(g)} \leftarrow$ Approx($A^{(g)}$, $\hat{p}^{(g)}$, $V^{(g)}$, K) and $Q', \sigma'_K, \Gamma' \leftarrow$ Approx(A', $\hat{p}'$, $V^{(g)}$, K)
if $\frac{\sigma_K^{(g)}}{\sqrt{|V^{(g)}|\hat{p}^{(g)}}} \cdot 1_{\{|V^{(g)}|\hat{p}^{(g)}\geq 50\}} \geq \frac{\sigma'_K}{\sqrt{|V^{(g)}|\hat{p}'}} \cdot 1_{\{|V^{(g)}|\hat{p}'\geq 50\}}$ **then**
  $(S_k)_{1\leq k\leq K} \leftarrow$ Detection ($Q^{(g)}, \Gamma^{(g)}, K$)
  Randomly place nodes in $V^{(g)} \setminus \Gamma^{(g)}$ to partitions $(S_k)_{k=1,\dots,K}$
else
  $(S_k)_{1\leq k\leq K} \leftarrow$ Detection (Q', Γ', K)
  Randomly place nodes in $V^{(g)} \setminus \Gamma'$ to partitions $(S_k)_{k=1,\dots,K}$
end if
Output: $(S_k)_{1\leq k\leq K}$,

Our algorithm to classify the green nodes consists in the following steps:
Step 1. Construct the indirect edge matrix A' using red nodes connected to two green nodes only.
Step 2. Perform a spectral analysis of matrices $A^{(g)}$ and A' as follows: first trim $A^{(g)}$ and A' (to remove nodes with too many connections), then extract their K largest eigenvalues and the corresponding eigenvectors.
Step 3. Select the matrix $A^{(g)}$ or A' with the largest normalized K-th largest eigenvalue.
Step 4. Construct the K clusters $V_1^{(g)}, \dots, V_K^{(g)}$ based on the eigenvectors of the matrix selected in the previous step.

The detailed pseudo-code of the algorithm is presented in Algorithm 1. Steps 2 and 4 of the algorithm are standard techniques used in clustering for the SBM, see e.g. [5]. The algorithms involved in these Steps are presented in the supplementary material (see Algorithms 4, 5, 6). Note that to extract the K largest eigenvalues and the corresponding eigenvectors of a matrix, we use the power method, which is memory-efficient (this becomes important when addressing Problem 2). Further observe that in Step 3, the algorithm exploits the information provided by the red nodes: it selects, between the direct edge matrix $A^{(g)}$ and the indirect edge matrix A', the matrix whose spectral properties provide more accurate information about the K clusters. This crucial step is enough for the algorithm to classify the green nodes asymptotically accurately whenever this is at all possible, as stated in the following theorem:

Theorem 4 *When $\sqrt{\gamma}f(n) = \omega(1)$, Algorithm 1 classifies the green nodes asymptotically accurately.*

In view of Theorem 2 (i), our algorithm is optimal. It might be surprising to choose one of the matrix $A^{(g)}$ or A' and throw the information contained in the other one. But the following simple calculation gives the main idea. To simplify, consider the case $\gamma f(n) = o(1)$ so that we know that the matrix $A^{(g)}$ alone is not sufficient to find the clusters. In this case, it is easy to see that the matrix A' alone allows to classify as soon as $\sqrt{\gamma}f(n) = \omega(1)$. Indeed, the probability of getting an indirect edge between two green nodes is of the order $(a^2+b^2)f(n)^2/(2n)$ if the two nodes are in the same clusters and $abf(n)^2/n$ if they are in different clusters. Moreover the graph of indirect edges has the same statistics as a SBM with these probabilities of connection. Hence standard results show that spectral methods will work as soon as $\gamma f(n)^2$ tends to infinity, i.e. the mean degree in the observed graph of indirect edges tends to infinity. In the case where $\gamma f(n)$ is too large (indeed $\gg \ln(f(n))$), then the graph of indirect edges becomes too sparse for A' to be useful. But in this regime, $A^{(g)}$ allows to classify the green nodes. This argument gives some intuitions about the full proof of Theorem 4 which can be found in the Appendix.

Algorithm 2 Greedy selections

Input: $A \in \{0,1\}^{|V|\times|V^{(g)}|}$, V, $V^{(g)}$, $(S_k^{(g)})_{1\le k\le K}$.
$V^{(r)} \leftarrow V \setminus V^{(g)}$ and $S_k \leftarrow S_k^{(g)}$, for all k
for $v \in V^{(r)}$ **do**
  Find $k^\star = \arg\max_k \{\sum_{w\in S_k^{(g)}} A_{vw}/|S_k^{(g)}|\}$ (tie broken uniformly at random)
  $S_{k^\star} \leftarrow S_{k^\star} \cup \{v\}$
end for
Output: $(S_k)_{1\le k\le K}$.

An attractive feature of our Algorithm 1 is that it does not require any parameter of the model as input except the number of clusters K. In particular, our algorithm selects automatically the best matrix among A' and $A^{(g)}$ based on their spectral properties.

Classifying red nodes. From Theorem 2 (ii), in order to classify red nodes, we need to assume that $\gamma f(n) = \omega(1)$. Under this assumption, the green nodes are well classified under Algorithm 1. To classify the red nodes accurately, we show that it is enough to greedily assign these nodes to the clusters of green nodes identified using Algorithm 1. More precisely, a red node v is assigned to the cluster that maximizes the number of observed edges between v and the green nodes of this cluster. The pseudo-code of this procedure is presented in Algorithm 2.

Theorem 5 *When $\gamma f(n) = \omega(1)$, combining Algorithms 1 and 2 yields an asymptotically accurate clustering algorithm.*

Again in view of Theorem 2 (ii), our algorithm is optimal. To summarize our results about Problem 1, i.e., clustering with partial information, we have shown that:
(a) If $\gamma \ll 1/f(n)^2$, no clustering algorithm can perform better than the naive algorithm that assigns nodes to clusters randomly (in the case of two clusters of equal sizes).
(b) If $1/f(n)^2 \ll \gamma \ll 1/f(n)$, Algorithm 1 classifies the green nodes asymptotically accurately, but no algorithm can classify the red nodes asymptotically accurately.
(c) If $1/f(n) \ll \gamma$, the combination of Algorithm 1 and Algorithm 2 classifies all nodes asymptotically accurately.

4 Clustering in the Streaming Model under Memory Constraints

In this section, we address Problem 2 where the clustering problem has additional constraints. Namely, the memory available to the algorithm is limited (memory constraints) and each column A_v of A is observed only once, hence if it is not stored, this information is lost (streaming model).

In view of previous results, when the entire matrix A is available (i.e. $\gamma = 1$) and when there is no memory constraint, we know that a necessary and sufficient condition for the existence of asymptotically accurate clustering algorithms is that $f(n) = \omega(1)$. Here we first devise a clustering algorithm adapted to the streaming model and using a memory scaling linearly with n that is asymptotically accurate as soon as $\log(n) \ll f(n)$. Algorithms 1 and 2 are the building blocks of this algorithm, and its performance analysis leverages the results of previous section. We also show that our algorithm does not need to sequentially observe all columns of A in order to accurately reconstruct the clusters. In other words, the algorithm uses strictly less than one pass on the data and is asymptotically accurate.

Clearly if the algorithm is asked (as above) to output the full partition of the network, it will require a memory scaling linearly with n, the size of the output. However, in the streaming model, we can remove this requirement and the algorithm can output the full partition sequentially similarly to an online algorithm (however our algorithm is not required to take an irrevocable action after the arrival of each column but will classify nodes after a group of columns arrives). In this case, the memory requirement can be sublinear. We present an algorithm with a memory requirement which depends on the density of the graph. In the particular case where $f(n) = n^\alpha$ with $0 < \alpha < 1$, our algorithm requires as little as n^β bits of memory with $\beta > \max\left(1-\alpha, \frac{2}{3}\right)$ to accurately cluster the nodes. Note that when the graph is very sparse ($\alpha \approx 0$), then the community detection is a hard statistical task and the algorithm needs to gather a lot of columns so that the memory requirement is quite

Algorithm 3 Streaming offline

Input: $\{A_1, \dots, A_T\}$, p, V, K

Initial: $N \leftarrow n \times K$ matrix filled with zeros and $B \leftarrow \frac{nh(n)}{\min\{np, n^{1/3}\}\log n}$

Subsampling: $A_t \leftarrow$ Randomly erase entries of A_t with probability $\max\{0, 1 - \frac{n^{1/3}}{np}\}$

for $\tau = 1$**to** $\lfloor \frac{T}{B} \rfloor$ **do**

  $A^{(B)} \leftarrow n \times B$ matrix where i-th column is $A_{i+(\tau-1)B}$

  $(S_k^{(\tau)}) \leftarrow$ Algorithm 1 $(A^{(B)}, V, \{(\tau-1)B+1, \dots, \tau B\}, K)$

  if $\tau = 1$ **then**

    $\hat{V}_k \leftarrow S_k^{(1)}$ for all k and $N_{v,k} \leftarrow \sum_{w \in S_k^{(1)}} A_{wv}$ for all $v \in V$ and k

  else

    $\hat{V}_{s(k)} \leftarrow \hat{V}_{s(k)} \cup S_k^{(\tau)}$ for all k where $s(k) = \arg\max_{1 \le i \le K} \frac{\sum_{v \in \hat{V}_i} \sum_{w \in S_k^{(\tau)}} A_{vw}}{|\hat{V}_i||S_k^{(\tau)}|}$

    $N_{v,s(k)} \leftarrow N_{v,s(k)} + \sum_{w \in S_k^{(\tau)}} A_{wv}$ for all $v \in V$ and k

  end if

end for

Greedy improvement : $\bar{V}_k \leftarrow \{v : k = \arg\max_{1 \le i \le K} \frac{N_{v,i}}{|\hat{V}_i|}\}$ for all k

Output: $(\bar{V}_k)_{1 \le k \le K}$,

high ($\beta \approx 1$). As α increases, the graph becomes denser and the statistical task easier. As a result, our algorithm needs to look at smaller blocks of columns and the memory requirement decreases. However, for $\alpha \geq 1/3$, although the statistical task is much easier, our algorithm hits its memory constraint and in order to store blocks with sufficiently many columns, it needs to subsample each column. As a result, the memory requirement of our algorithm does not decrease for $\alpha \geq 1/3$.

The main idea of our algorithms is to successively treat *blocks* of B consecutive arriving columns. Each column of a block is stored in the memory. After the last column of a block arrives, we apply Algorithm 1 to classify the corresponding nodes accurately, and we then merge the obtained clusters with the previously identified clusters. In the online version, the algorithm can output the partition of the block and in the offline version, it stores this result. We finally remove the stored columns, and proceed with the next block. For the offline algorithm, after a total of T observed columns, we apply Algorithm 2 to classify the remaining nodes so that T can be less than n. The pseudo-code of the offline algorithm is presented in Algorithm 3. Next we discuss how to tune B and T so that the classification is asymptotically accurate, and we compute the required memory to implement the algorithm.

Block size. We denote by B the size of a block. Let $h(n)$ be such that the block size is $B = \frac{h(n)n}{f(n)\log(n)}$. Let $\bar{f}(n) = \min\{f(n), n^{1/3}\}$ which represents the order of the number of positive entries of each column after the subsampling process. According to Theorem 4 (with $\gamma = B/n$), to accurately classify the nodes arrived in a block, we just need that $\frac{B}{n}\bar{f}(n)^2 = \omega(1)$, which is equivalent to $h(n) = \omega(\frac{\log(n)}{\min\{f(n), n^{1/3}\}})$. Now the merging procedure that combines the clusters found analyzing the current block with the previously identified clusters uses the number of connections between the nodes corresponding to the columns of the current block to the previous clusters. The number of these connections must grow large as n tends to ∞ to ensure the accuracy of the merging procedure. Since the number of these connections scales as $B^2\frac{\bar{f}(n)}{n}$, we need that $h(n)^2 = \omega(\min\{f(n), n^{1/3}\}\frac{\log(n)^2}{n})$. Note that this condition is satisfied as long as $h(n) = \omega(\frac{\log(n)}{\min\{f(n), n^{1/3}\}})$.

Total number of columns for the offline algorithm. To accurately classify the nodes whose columns are not observed, we will show that we need the total number of observed columns T to satisfy $T = \omega(\frac{n}{\min\{f(n), n^{1/3}\}})$ (which is in agreement with Theorem 5).

Required memory for the offline algorithm. To store the columns of a block, we need $\Theta(nh(n))$ bits. To store the previously identified clusters, we need at most $\log_2(K)n$ bits, and we can store the number of connections between the nodes corresponding to the columns of the current block to the previous clusters using a memory linearly scaling with n. Finally, to execute Algorithm 1, the

Algorithm 4 Streaming online

Input: $\{A_1, \ldots, A_n\}$, p, V, K
Initial: $B \leftarrow \frac{nh(n)}{\min\{np, n^{1/3}\}\log n}$ and $\tau^\star = \lfloor \frac{T}{B} \rfloor$
Subsampling: $A_t \leftarrow$ Randomly erase entries of A_t with probability $\max\{0, 1 - \frac{n^{1/3}}{np}\}$
for $\tau = 1$**to** $\tau^\star$ **do**
  $A^{(B)} \leftarrow n \times B$ matrix where i-th column is $A_{i+(\tau-1)B}$
  $(S_k)_{1\le k\le K} \leftarrow$ Algorithm 1 $(A^{(B)}, V, \{(\tau-1)B+1, \ldots, \tau B\}, K)$
  if $\tau = 1$ **then**
    $\hat{V}_k \leftarrow S_k$ for all k
    Output at B**:** $(S_k)_{1\le k\le K}$
  else
    $s(k) \leftarrow \arg\max_{1\le i\le K} \frac{\sum_{v\in\hat{V}_i}\sum_{w\in S_k} A_{vw}}{|\hat{V}_i||S_k|}$ for all k
    Output at τB**:** $(S_{s(k)})_{1\le k\le K}$
  end if
end for

power method used to perform the SVD (see Algorithm 5) requires the same amount of bits than that used to store a block of size B. In summary, the required memory is $M = \Theta(nh(n) + n)$.

Theorem 6 *Assume that* $h(n) = \omega(\frac{\log(n)}{\min\{f(n), n^{1/3}\}})$ *and* $T = \omega(\frac{n}{\min\{f(n), n^{1/3}\}})$. *Then with* $M = \Theta(nh(n) + n)$ *bits, Algorithm 3, with block size* $B = \frac{h(n)n}{\min\{f(n), n^{1/3}\}\log(n)}$ *and acquiring the* T *first columns of* A*, outputs clusters* $\hat{V}_1, \ldots, \hat{V}_K$ *such that with high probability, there exists a permutation* σ *of* $\{1, \ldots, K\}$ *such that:* $\frac{1}{n}\left|\bigcup_{1\le k\le K} \hat{V}_k \setminus V_{\sigma(k)}\right| = O\left(\exp(-cT\frac{\min\{f(n), n^{1/3}\}}{n})\right)$ *with a constant* $c > 0$.

Under the conditions of the above theorem, Algorithm 3 is asymptotically accurate. Now if $f(n) = \omega(\log(n))$, we can choose $h(n) = 1$. Then Algorithm 3 classifies nodes accurately and uses a memory linearly scaling with n. Note that increasing the number of observed columns T just reduces the proportion of misclassified nodes. For example, if $f(n) = \log(n)^2$, with high probability, the proportion of misclassified nodes decays faster than $1/n$ if we acquire only $T = n/\log(n)$ columns, whereas it decays faster than $\exp(-\log(n)^2)$ if all columns are observed.

Our online algorithm is a slight variation of the offline algorithm. Indeed, it deals with the first block exactly in the same manner and keeps in memory the partition of this first block. It then handles the successive blocks as the first block and merges the partition of these blocks with those of the first block as done in the offline algorithm for the second block. Once this is done, the online algorithm just throw all the information away except the partition of the first block.

Theorem 7 *Assume that* $h(n) = \omega(\frac{\log(n)}{\min\{f(n), n^{1/3}\}})$*, then Algorithm 4 with block size* $B = \frac{h(n)n}{\min\{f(n), n^{1/3}\}\log n}$ *is asymptotically accurate (i.e., after one pass, the fraction of misclassified nodes vanishes) and requires* $\Theta(nh(n))$ *bits of memory.*

5 Conclusion

We introduced the problem of community detection with partial information, where only an induced subgraph corresponding to a fraction of the nodes is observed. In this setting, we gave a necessary condition for accurate reconstruction and developed a new spectral algorithm which extracts the clusters whenever this is at all possible. Building on this result, we considered the streaming, memory limited problem of community detection and developed algorithms able to asymptotically reconstruct the clusters with a memory requirement which is linear in the size of the network for the offline version of the algorithm and which is sublinear for its online version. To the best of our knowledge, these algorithms are the first community detection algorithms in the data stream model. The memory requirement of these algorithms is non-increasing in the density of the graph and determining the optimal memory requirement is an interesting open problem.

References

[1] R. B. Boppana. Eigenvalues and graph bisection: An average-case analysis. In *Foundations of Computer Science, 1987., 28th Annual Symposium on*, pages 280–285. IEEE, 1987.

[2] S. Chatterjee. Matrix estimation by universal singular value thresholding. *arXiv preprint arXiv:1212.1247*, 2012.

[3] K. Chaudhuri, F. C. Graham, and A. Tsiatas. Spectral clustering of graphs with general degrees in the extended planted partition model. *Journal of Machine Learning Research-Proceedings Track*, 23:35–1, 2012.

[4] Y. Chen, S. Sanghavi, and H. Xu. Clustering sparse graphs. In *Advances in Neural Information Processing Systems 25*, pages 2213–2221. 2012.

[5] A. Coja-Oghlan. Graph partitioning via adaptive spectral techniques. *Combinatorics, Probability & Computing*, 19(2):227–284, 2010.

[6] A. Dasgupta, J. Hopcroft, R. Kannan, and P. Mitra. Spectral clustering by recursive partitioning. In *Algorithms–ESA 2006*, pages 256–267. Springer, 2006.

[7] A. Decelle, F. Krzakala, C. Moore, and L. Zdeborová. Inference and phase transitions in the detection of modules in sparse networks. *Phys. Rev. Lett.*, 107, Aug 2011.

[8] U. Feige and E. Ofek. Spectral techniques applied to sparse random graphs. *Random Structures & Algorithms*, 27(2):251–275, 2005.

[9] S. Guha, N. Mishra, R. Motwani, and L. O'Callaghan. Clustering data streams. In *41st Annual Symposium on Foundations of Computer Science (Redondo Beach, CA, 2000)*, pages 359–366. IEEE Comput. Soc. Press, Los Alamitos, CA, 2000.

[10] P. Holland, K. Laskey, and S. Leinhardt. Stochastic blockmodels: First steps. *Social Networks*, 5(2):109 – 137, 1983.

[11] M. Jerrum and G. B. Sorkin. The metropolis algorithm for graph bisection. *Discrete Applied Mathematics*, 82(13):155 – 175, 1998.

[12] L. Massoulié. Community detection thresholds and the weak ramanujan property. *CoRR*, abs/1311.3085, 2013.

[13] F. McSherry. Spectral partitioning of random graphs. In *Foundations of Computer Science, 2001. Proceedings. 42nd IEEE Symposium on*, pages 529–537. IEEE, 2001.

[14] I. Mitliagkas, C. Caramanis, and P. Jain. Memory limited, streaming PCA. In *NIPS*, 2013.

[15] E. Mossel, J. Neeman, and A. Sly. Stochastic block models and reconstruction. *arXiv preprint arXiv:1202.1499*, 2012.

[16] M. Rudelson and R. Vershynin. The littlewood–offord problem and invertibility of random matrices. *Advances in Mathematics*, 218(2):600–633, 2008.

[17] S. Yun and A. Proutiere. Community detection via random and adaptive sampling. In *COLT*, 2014.

A Algorithms

We present below three algorithms that constitute building blocks of the main algorithms presented in the paper.

Algorithm 5 Approx (A, $\hat{p}$, V, K)

Input: A, $\hat{p}$, V, K
$\ell^\star \leftarrow \max\{1, \lfloor |V| \exp(-|V|\hat{p}) \rfloor\})$
for $v \in V$ **do**
  $x_v \leftarrow \sum_{w \in V} A_{vw}$
end for
$x^\star \leftarrow \ell^\star$-th largest x_v
$\Gamma \leftarrow \{v | x_v \leq x^\star,\ v \in V\}$
$\bar{A} \leftarrow (A_{vw})_{v,w \in \Gamma}$
$(Q, \sigma_K) \leftarrow$ Power Method $(\bar{A}, \Gamma, K)$ (Algorithm 6)
Output: (Q, σ_K, Γ)

Algorithm 6 Power Method (A, V, K)

Input: A, V, K
Initial: $Q_0 \leftarrow$ Randomly choose K orthonormal vectors and $\tau^\star = \lceil \log |V| \rceil$
for $\tau = 1$ **to** $\tau^\star$ **do**
  $AQ_{\tau-1} = Q_\tau R_\tau$
end for
$\sigma_K \leftarrow K$-th largest singular value of $R_{\tau^\star}$
Output: $(Q_{\tau^\star}, \sigma_K)$

Algorithm 7 Detection (Q, V, K)

Input: Q, V, K (let Q_v denote the low of Q corresponding to v)
for $i = 1$ **to** $\log |V|$ **do**
  $X_{i,v} \leftarrow \{w \in V : \|Q_w - Q_v\|^2 \leq \frac{i}{|V| \log |V|}\}$
  $T_{i,0} \leftarrow \emptyset$
  for $k = 1$ **to** K **do**
    $v_k^\star \leftarrow \arg\max_v |X_{i,v} \setminus \bigcup_{l=1}^{k-1} T_{i,l}|$
    $T_{i,k} \leftarrow X_{i,v_k^\star} \setminus \bigcup_{l=1}^{k-1} T_{i,l}$ and $\xi_{i,k} \leftarrow \sum_{v \in T_{i,k}} Q_v / |T_{i,k}|$.
  end for
  for $v \in V \setminus (\bigcup_{k=1}^{K} T_{i,k})$ **do**
    $k^\star \leftarrow \arg\min_k \|Q_v - \xi_{i,k}\|$
    $T_{i,k^\star} \leftarrow T_{i,k^\star} \cup \{v\}$
  end for
  $r_i \leftarrow \sum_{k=1}^{K} \sum_{v \in T_{i,k}} \|Q_v - \xi_{i,k}\|^2$
end for
$i^\star \leftarrow \arg\min_i r_i$.
$S_k \leftarrow T_{i^\star,k}$ for all k
Output: $(S_k)_{k=1,\ldots,K}$.

B Proofs

In the following, we denote by $\lambda_i(X)$ the i-th largest singular value of matrix X.

B.1 Proof of Theorem 1

Preliminaries. In what follows, we denote by $\sigma^{(g)} \in \{-1, 1\}^{n^{(g)}}$ a vector that represents the repartition of nodes in the two communities, i.e., nodes v and w belong to the same community if and only if $\sigma_v^{(g)} = \sigma_w^{(g)}$. We also denote by $\hat{\sigma}^{(g)} \in \{-1, 1\}^{n^{(g)}}$ the estimate of $\sigma^{(g)}$ that a clustering algorithm could return.

We further introduce the following notation. For any $k > 0$ and any two vectors $x, y \in \{-1, 1\}^k$, we denote by $d_H(x, y) = \sum_{i=1}^k 1(x_i \neq y_i)$ the Hamming distance between x and y and define

$$d(x, y) = \frac{1}{k} \min\{d_H(x, y), d_H(x, -y)\}.$$

For an estimate $\hat{\sigma}^{(g)}$ of $\sigma^{(g)}$, the quantity $d(\hat{\sigma}^{(g)}, \sigma^{(g)})$ is exactly the fraction of misclassified green nodes. Hence if estimate $\hat{\sigma}^{(g)}$ is obtained from algorithm π, we have $\epsilon^\pi(n^{(g)}) = d(\hat{\sigma}^{(g)}, \sigma^{(g)})$. Note that $d(\hat{\sigma}^{(g)}, \sigma^{(g)}) \leq 1/2$.

We first state key lemmas for this proof. Their proofs are postponed to the end of this section.

Lemma 8 *For any $\alpha < 1/2$ and estimate $\hat{\sigma}^{(g)}$, we have as $n^{(g)} \to \infty$*

$$\mathbb{P}(d(\hat{\sigma}^{(g)}, \sigma^{(g)}) > \alpha) \geq 1 - \frac{n^{(g)} - H(\sigma^{(g)}|A)}{n^{(g)}(1 - H(\alpha))} + o(1),$$

where $H(\alpha) = -\alpha \log \alpha - (1 - \alpha) \log(1 - \alpha)$ and $H(\sigma^{(g)}|A)$ is the conditional entropy of $\sigma^{(g)}$ knowing A.

Lemma 9 *As $n^{(g)} \to \infty$, we have:*

$$H(A) - H(A|\sigma^{(g)}) \leq o(n^{(g)}) + O(n^{(g)} \gamma f(n)^2).$$

From the definition of conditional entropy, we have

$$H(\sigma^{(g)}|A) = H(\sigma^{(g)}) - H(A) + H(A|\sigma^{(g)}) = n^{(g)}(1 - o(1)), \tag{1}$$

since $H(\sigma^{(g)}) = \log \binom{n^{(g)}}{n^{(g)}/2} \geq n^{(g)} - \frac{1}{2} \log 2n^{(g)}$ and we have $H(A) - H(A|\sigma^{(g)}) = o(n^{(g)})$ from Lemma 9. As soon as $n^{(g)} \to \infty$, putting (1) into Lemma 8, we see that for any $\alpha < 1/2$ and any estimate $\hat{\sigma}^{(g)}$,

$$\mathbb{P}(d(\hat{\sigma}^{(g)}, \sigma^{(g)}) > \alpha) \to 1.$$

If $\hat{\sigma}^{(g)}$ is a random guess, i.e. for each $v \in V^{(g)}$, $\hat{\sigma}_v^{(g)}$ is equal to 1 or -1 with probability $1/2$ independently of the rest, then for any $\alpha < 1/2$, as soon as $n^{(g)} \to \infty$, we have by the weak law of large numbers, $\mathbb{P}(d(\hat{\sigma}^{(g)}, \sigma^{(g)}) > \alpha) \to 1$. Since we have

$$\mathbb{E}[\epsilon^\pi(n^{(g)})] \geq \alpha \mathbb{P}(d(\hat{\sigma}^{(g)}, \sigma^{(g)}) > \alpha),$$

and α can be chosen as close to $1/2$ as desired, the result follows.

B.2 Proof of Lemma 8

We define the event $E = \{d(\hat{\sigma}^{(g)}, \sigma^{(g)}) > \alpha\}$ and P_e its probability. We have

$$\begin{aligned} H(E, \sigma^{(g)}|\hat{\sigma}^{(g)}) &= H(\sigma^{(g)}|\hat{\sigma}^{(g)}) + \underbrace{H(E|\sigma^{(g)}, \hat{\sigma}^{(g)})}_{0} \\ &= H(E|\hat{\sigma}^{(g)}) + H(\sigma^{(g)}|E, \hat{\sigma}^{(g)}) \\ &\leq H(P_e) + P_e \log \binom{n^{(g)}}{n^{(g)}/2} + (1 - P_e)(n^{(g)} H(\alpha) + \log n^{(g)}), \end{aligned}$$

where the last inequality follows from $H(E|\hat{\sigma}^{(g)}) \leq H(E) = H(P_e)$ and the fact that

$$|\{\sigma^{(g)}, d(\sigma^{(g)}, \hat{\sigma}^{(g)}) \leq \alpha\}| = \sum_{i=0}^{\alpha n^{(g)}} \binom{n^{(g)}}{i} \leq (n^{(g)}\alpha + 1) \binom{n^{(g)}}{\alpha n^{(g)}} \leq n^{(g)} 2^{n^{(g)} H(\alpha)}.$$

Using $H(P_e) \leq 1$ and that $\binom{n^{(g)}}{n^{(g)}/2} \leq 2^{n^{(g)}}$ for sufficiently large $n^{(g)}$, we get

$$P_e \geq \frac{H(\sigma^{(g)}|\hat{\sigma}^{(g)}) - 1 - n^{(g)} H(\alpha) - \log n^{(g)}}{n^{(g)}(1 - H(\alpha)) - \log n^{(g)}}.$$

The claim follows from the data processing inequality which ensures $H(\sigma^{(g)}|\hat{\sigma}^{(g)}) \geq H(\sigma^{(g)}|A)$.

B.3 Proof of Lemma 9

Thanks to independence, we have

$$H(A) - H(A|\sigma^{(g)}) \quad = \quad H(A^{(g)}) - H(A^{(g)}|\sigma^{(g)}) + H(A^{(r)}) - H(A^{(r)}|\sigma^{(g)})$$

We first deal with the first term $H(A^{(g)}) - H(A^{(g)}|\sigma^{(g)})$. By the concavity of $p \mapsto H(p)$, we have $H(A^{(g)}) \leq \binom{n^{(g)}}{2} H\left(\frac{p+q}{2}\right)$ and $H(A^{(g)}|\sigma^{(g)}) = 2\binom{n^{(g)}/2}{2} H(p) + (\frac{n^{(g)}}{2})^2 H(q)$. Hence, we get

$$\begin{aligned}
H(A^{(g)}) - H(A^{(g)}|\sigma^{(g)}) \leq & \binom{n^{(g)}}{2} H(\frac{p+q}{2}) - 2\binom{n^{(g)}/2}{2} H(p) - (\frac{n^{(g)}}{2})^2 H(q) \\
= & \frac{(n^{(g)})^2}{4} \left(2H(\frac{p+q}{2}) - H(p) - H(q) \right) + o(n^{(g)}) \\
= & \frac{(n^{(g)})^2}{4} \left(p \log \frac{2p}{p+q} + q \log \frac{2q}{p+q} \right) \\
& + \frac{(n^{(g)})^2}{4} \left((1-p) \log \frac{2-2p}{2-p-q} + (1-q) \log \frac{2-2q}{2-p-q} \right) + o(n^{(g)}) \\
\leq & \frac{(n^{(g)})^2}{4} \left(\frac{(p-q)^2}{p+q} + \frac{(p-q)^2}{2-p-q} \right) + o(n^{(g)}) \\
= & o(n^{(g)}) + o(n^{(g)} \gamma f(n)).
\end{aligned}$$

We denote by $A_v^{(r)}$ the row vector of $A^{(r)}$ corresponding to $v \in V^{(r)}$. For the second term, by independence we have

$$H(A^{(r)}) - H(A^{(r)}|\sigma^{(g)}) = (n - n^{(g)}) \big(H(A_v^{(r)}) - H(A_v^{(r)}|\sigma^{(g)}) \big).$$

For a vector $x \in \{-1,1\}^{V^{(g)}}$ and $\sigma^{(g)}$, we define $|x|^+ = \sum_{v \in V^{(g)},\, \sigma_v^{(g)}=1} x_v$, $|x|^- = \sum_{v \in V^{(g)},\, \sigma_v^{(g)}=-1} x_v$ and $|x| = |x|^+ + |x|^-$. For a given $\sigma^{(g)}$, we have

$$\mathbb{P}[A_v^{(r)} = x | \sigma^{(g)}] = \zeta(|x|^+, |x|^-),$$

where

$$\zeta(i,j) = \frac{((\frac{p}{1-p})^i (\frac{q}{1-q})^j + (\frac{p}{1-p})^j (\frac{q}{1-q})^i)(1-p)^{\frac{m}{2}} (1-q)^{\frac{m}{2}}}{2}.$$

Since $\sigma^{(g)}$ is uniformly distributed,

$$\begin{aligned}
\mathbb{P}[A_v^{(g)} = x] = & \binom{n^{(g)}}{n^{(g)}/2}^{-1} \sum_{\sigma^{(g)} : \sum_{v \in V^{(g)}} \sigma_v^{(g)} = 0} \mathbb{P}[A_v^{(g)} = x | \sigma^{(g)}] \\
= & \binom{n^{(g)}}{n^{(g)}/2}^{-1} \sum_{i=0}^{|x|} \binom{|x|}{i} \binom{n^{(g)} - |x|}{n^{(g)}/2 - i} \zeta(i, |x| - i) = \eta(|x|),
\end{aligned}$$

where

$$\eta(k) = \frac{\sum_{i=0}^{k} \binom{n^{(g)}/2}{i} \binom{n^{(g)}/2}{k-i} \zeta(i, k-i)}{\sum_{i=0}^{k} \binom{n^{(g)}/2}{i} \binom{n^{(g)}/2}{k-i}}.$$

Since $\eta(0)=\zeta(0,0)$, $\eta(1)=\zeta(1,0)=\zeta(0,1)$, and $\binom{n^{(g)}}{k}\eta(k)=\sum_{i=0}^{k}\binom{n^{(g)}/2}{i}\binom{n^{(g)}/2}{k-i}\zeta(i,k-i)$,

$$
\begin{aligned}
&H(A_v^{(r)})-H(A_v^{(r)}|\sigma^{(g)})\\
&=-\sum_{k=0}^{n^{(g)}}\binom{n^{(g)}}{k}\eta(k)\log\eta(k)+\sum_{i=0}^{n^{(g)}/2}\sum_{j=0}^{n^{(g)}/2}\binom{n^{(g)}/2}{i}\binom{n^{(g)}/2}{j}\zeta(i,j)\log\zeta(i,j)\\
&=\sum_{i=0}^{n^{(g)}/2}\sum_{j=0}^{n^{(g)}/2}\binom{n^{(g)}/2}{i}\binom{n^{(g)}/2}{j}\zeta(i,j)\log\frac{\zeta(i,j)}{\eta(i+j)}\\
&=\sum_{k=2}^{n^{(g)}}\sum_{i=0}^{k}1_{i\le n^{(g)}/2}1_{k-i\le n^{(g)}/2}\binom{n^{(g)}/2}{i}\binom{n^{(g)}/2}{k-i}\zeta(i,k-i)\log\frac{\zeta(i,k-i)}{\eta(k)}\\
&\le\sum_{k=2}^{n^{(g)}}\sum_{i=0}^{k}1_{i\le n^{(g)}/2}1_{k-i\le n^{(g)}/2}\binom{n^{(g)}/2}{i}\binom{n^{(g)}/2}{k-i}\zeta(i,k-i)k\log(\frac{p}{q})\\
&\le\sum_{2\le k\le n^{(g)}}(n^{(g)}p)^k k\log(\frac{a}{b})\\
&=O((n^{(g)})^2p^2),
\end{aligned}
$$

where the last equality stems from $n^{(g)}p=o(1)$. Thus,

$$(n-m)\big(H(A_v^{(r)})-H(A_v^{(r)}|\sigma^{(g)})\big)=O(n(n^{(g)})^2p^2)=O(n^{(g)}\frac{n^{(g)}}{n}f(n)^2)=O(n^{(g)}\gamma f(n)^2),$$

and the lemma follows since $f(n)\ge 1$ so that $f(n)^2\ge f(n)$.

B.4 Proof of Theorem 2

In the remaining proofs, we use m instead of $n^{(g)}$ to denote the number of green nodes. We first consider case (i) with $\gamma=\Theta(1)$. In this case, a necessary condition for the existence of an asymptotically accurate clustering algorithm is that the fraction of green nodes outside the largest connected component of the observed graph vanishes as $n\to\infty$. This condition imposes that $f(n)\to\infty$.

We now consider case (i) with $m=o(n)$, i.e. $\gamma=o(1)$. Denote by Φ the true hidden partition $(V_1^{(g)},V_2^{(g)})$ for green nodes. Let $\mathbb{P}_\Phi$ be the probability measure capturing the randomness in the observations assuming that the network structure is described by Φ. We also introduce a slightly different structure Ψ. The latter is described by clusters $V_1'^{(g)}=V_1^{(g)}\cup\{v_2\}\setminus\{v_1\}$, $V_2'^{(g)}=V_2^{(g)}\cup\{v_1\}\setminus\{v_2\}$ with arbitrary selected $v_1\in V_1^{(g)}$ and $v_2\in V_2^{(g)}$.

Let $\pi\in\Pi$ denote a clustering algorithm for green nodes with output $(\hat{V}_1^{(g)},\hat{V}_2^{(g)})$, and let $\mathcal{E}=\hat{V}_1^{(g)}\triangle V_1^{(g)}$ be the set of misclassified nodes under π. Note that in general in our proofs, we always assume without loss of generality that $|\hat{V}_1^{(g)}\triangle V_1^{(g)}|\le|\hat{V}_1^{(g)}\triangle V_2^{(g)}|$, so that the set of misclassified nodes is really $\mathcal{E}$. Further define $\mathcal{B}=\{v_1\in\hat{V}_1^{(g)}\}$ as the set of events where node v_1 is correctly classified. We have $\varepsilon(m)=|\mathcal{E}|/m$.

Let $x_{i,j}$ be equal to one if there is an edge between nodes i and j and zero otherwise.With a slight abuse of notation, we define the boolean functions $p(\cdot)$ and $q(\cdot)$ as follows: $p(1)=af(n)/n=p$, $q(1)=bf(n)/n=q$ and $p(0)=1-p(1)$, $q(0)=1-q(1)$. We introduce L (a quantity that resembles the log-likelihood ratio between $\mathbb{P}_\Phi$ and $\mathbb{P}_\Psi$) as:

$$
\begin{aligned}
L=&\sum_{i\in V_1'^{(g)}}\log\frac{q(x_{i,v_1})p(x_{i,v_2})}{p(x_{i,v_1})q(x_{i,v_2})}+\sum_{i\in V_2'^{(g)}}\log\frac{p(x_{i,v_1})q(x_{i,v_2})}{q(x_{i,v_1})p(x_{i,v_2})}\\
&+\sum_{v\in V^{(r)}}\log\frac{\prod_{i\in V_1'^{(g)}}p(x_{v,i})\prod_{i\in V_2'^{(g)}}q(x_{v,i})+\prod_{i\in V_1'^{(g)}}q(x_{v,i})\prod_{i\in V_2'^{(g)}}p(x_{v,i})}{\prod_{i\in V_1^{(g)}}p(x_{v,i})\prod_{i\in V_2^{(g)}}q(x_{v,i})+\prod_{i\in V_1^{(g)}}q(x_{v,i})\prod_{i\in V_2^{(g)}}p(x_{v,i})},
\end{aligned}
$$

In what follows, we establish a relationship between $\mathbb{E}[\varepsilon(m)]$ and L. For any function $g(m)$,

$$\mathbb{P}_\Psi\{L \le g(m)\} \quad = \quad \mathbb{P}_\Psi\{L \le g(m), \bar{\mathcal{B}}\} + \mathbb{P}_\Psi\{L \le g(m), \mathcal{B}\}. \tag{2}$$

We have:

$$\begin{aligned}
\mathbb{P}_\Psi\{L \le g(m), \bar{\mathcal{B}}\} &= \int_{\{L \le g(m), \bar{\mathcal{B}}\}} d\mathbb{P}_\Psi \\
&= \int_{\{L \le g(m), \bar{\mathcal{B}}\}} \prod_{i \in V_1'} \frac{\nu(x_{i,v_1})\nu(x_{i,v_2})}{p(x_{i,v_1})q(x_{i,v_2})} \prod_{i \in V_2'} \frac{\nu(x_{i,v_1})\nu(x_{i,v_2})}{q(x_{i,v_1})p(x_{i,v_2})} d\mathbb{P}_\Phi \\
&\le \exp(g(m))\mathbb{P}_\Phi\{L \le g(m), \bar{\mathcal{B}}\} \le \exp(g(m))\mathbb{P}_\Phi\{\bar{\mathcal{B}}\} \\
&\le 2\exp(g(m))\mathbb{E}_\Phi[\varepsilon(m)],
\end{aligned} \tag{3}$$

where the last inequality comes from the fact that,

$$\mathbb{P}_\Phi\{\mathcal{B}\} \quad \ge \quad 1 - \mathbb{P}_\Phi\{v_1 \notin \hat{V}_1^{(g)}\} \ge 1 - 2\mathbb{E}_\Phi[\varepsilon(n)].$$

We also have:

$$\mathbb{P}_\Psi\{L \le g(m), \mathcal{B}\} \le \mathbb{P}_\Psi\{\mathcal{B}\} \le 2\mathbb{E}_\Psi[\varepsilon(m)]. \tag{4}$$

By (3) and (4)

$$\mathbb{P}_\Psi\{L \le g(n)\} \le 2\mathbb{E}_\Phi[\varepsilon(n)]\exp(g(n)) + 2\mathbb{E}_\Psi[\varepsilon(n)].$$

Since $\mathbb{E}_\Phi[\varepsilon(n)] = \mathbb{E}_\Psi[\varepsilon(n)] = \mathbb{E}[\varepsilon(n)]$ and $\mathbb{E}[\varepsilon(n)] = o(1)$, choosing $g(m) = \log\left(\frac{1}{8\mathbb{E}[\varepsilon(m)]}\right)$, we obtain:

$$\lim\inf_{m\to\infty} \mathbb{P}_\Psi\{L \ge \log\left(\frac{1}{8\mathbb{E}[\varepsilon(m)]}\right)\} > \frac{1}{2}. \tag{5}$$

By Chebyshev's inequality, $\mathbb{P}_\Psi\{L \ge \mathbb{E}_\Psi[L] + 2\sigma_\Psi[L]\} \le \frac{1}{4}$. Therefore, to be valid the above inequality,

$$\mathbb{E}_\Psi[L] + 2\sigma_\Psi[L] \ge \log\left(\frac{1}{8\mathbb{E}[\varepsilon(m)]}\right), \tag{6}$$

which implies that $\mathbb{E}_\Psi[L] + 2\sigma_\Psi[L] = \omega(1)$ since $\mathbb{E}[\varepsilon(m)] = o(1)$.

We define $KL(p,q) = p\log(p/q) + (1-p)\log((1-p)/(1-q))$. From the definition of L, we can easily bound $\mathbb{E}_\Psi[L]$ and $\sigma_\Psi[L]^2$:

$$\begin{aligned}
\mathbb{E}_\Psi[L] \le& m \cdot (KL(p,q) + KL(q,p)) \\
&+ n \sum_{0 \le i,j \le \frac{m}{2}-1} \binom{m/2-1}{i}\binom{m/2-1}{j} \frac{p^{i+1}q^j + p^j q^{i+1}}{2} \log \frac{p^{i+1}q^j + p^j q^{i+1}}{p^i q^{j+1} + p^{j+1}q^i} \\
\le& m \cdot (KL(p,q) + KL(q,p)) + n \sum_{1 \le k \le m} m^k p^{k+1} k \log \frac{p}{q} \\
\le& O(\gamma f(n)) + np \log \frac{a}{b} \sum_{1 \le k \le m} k m^k p^k \le O(\gamma f(n)) + np \log \frac{a}{b} \sum_{k=1}^{\infty} k(mp)^k \\
\le& O(\gamma f(n)) + np \log \frac{a}{b} \sum_{k=1}^{\infty} (2mp)^k \\
\sigma_\Psi[L]^2 \le& m((p+q)(\log \frac{a}{b})^2 + (2-p-q)(\log \frac{1-q}{1-p})^2) \\
&+ n \sum_{0 \le i,j \le \frac{m}{2}-1} \binom{m/2-1}{i}\binom{m/2-1}{j} \frac{p^{i+1}q^j + p^j q^{i+1}}{2} \left(\log \frac{p^{i+1}q^j + p^j q^{i+1}}{p^i q^{j+1} + p^{j+1}q^i}\right)^2 \\
\le& 4mp(\log \frac{a}{b})^2 + n \sum_{1 \le k \le m} m^k p^{k+1} k^2 (\log \frac{a}{b})^2 \\
\le& O(\gamma f(n)) + np(\log \frac{a}{b})^2 \sum_{k=1}^{\infty} k^2 (mp)^k \le O(\gamma f(n)) + np(\log \frac{a}{b})^2 \sum_{k=1}^{\infty} (3mp)^k.
\end{aligned} \tag{7}$$

Therefore, the necessary condition for $\mathbb{E}_\Psi[L] + 2\sigma_\Psi[L] = \omega(1)$ is that $np\sum_{k=1}^{\infty}(3mp)^k = \omega(1)$. We conclude this proof from that $np\sum_{k=1}^{\infty}(3mp)^k = \omega(1)$ if and only if $\gamma f(n)^2 = \omega(1)$.

We now prove point (ii). Note that the probability for a red node to be isolated is at least $(1 - af(n)/n)^{\gamma n} \approx \exp(-a\gamma f(n))$. If there exists an asymptotically accurate clustering algorithm, then the fraction of such isolated red nodes should vanishes and hence $\gamma f(n) \to \infty$.

B.5 Proof of Theorem 4

The proof proceeds in two steps. Step 1. We first establish that if $\frac{\sigma_K^{(g)}}{\sqrt{m\hat{p}^{(g)}}} \cdot 1_{\{m\hat{p}^{(g)} \geq 50\}} = \omega(1)$, then the spectral method applied to the matrix $A^{(g)}$ is asymptotically accurate. We also show that if $\frac{\sigma'_K}{\sqrt{m\hat{p}'}} \cdot 1_{\{m\hat{p}' \geq 50\}} = \omega(1)$, then the spectral method applied to the matrix of indirect edges A' is asymptotically accurate. Step 2. We show that if $\gamma f(n) = \omega(1)$, then $\frac{\sigma_K^{(g)}}{\sqrt{m\hat{p}^{(g)}}} \cdot 1_{\{m\hat{p}^{(g)} \geq 50\}} = \omega(1)$ with high probability (w.h.p.), and if $\gamma f(n) = O(1)$ and $\sqrt{\gamma} f(n) = \omega(1)$, then $\frac{\sigma'_K}{\sqrt{m\hat{p}'}} \cdot 1_{\{m\hat{p}' \geq 50\}} = \omega(1)$ w.h.p..

Preliminaries. We first state three lemmas to analyze the performance of Approx, PowerMethod, and Detection algorithms. Their proofs are postponed to the end of this section. In what follows, let $V = \{1, \ldots, n\}$ and let $A \in \mathbb{R}^{n \times n}$. For any matrix $Z \in \mathbb{R}^{n \times n}$, $\lambda_K(Z)$ denotes the K-th largest singular value of Z.

Lemma 10 *With probability $1 - O(1/n)$, the output (Q, σ_K) of the PowerMethod algorithm with input (A, V, K) (Algorithm 6) satisfies that $\sigma_K = \Theta(\lambda_K(A))$.*

Lemma 11 *Let $A, M \in \mathbb{R}^{n \times n}$ and let $M = U\Lambda U^T$ be the SVD of M where $\Lambda \in \mathbb{R}^{K \times K}$. Assume that $\|A - M\| = o(\lambda_K(M))$, the output (Q, σ_K) of the PowerMethod algorithm (Algorithm 6) with input (A, V, K) satisfies:*

$$\|U_\perp^T Q\| = O\left(\frac{\|A - M\|}{\lambda_K(M)}\right) = o(1),$$

where $U_\perp$ is an orthonormal basis of the space perpendicular to the linear span of U.

Lemma 12 *Assume that the set V is partitioned into K subsets $(V_k)_{1 \leq k \leq K}$. Further assume that for any k, $\frac{|V_k|}{n} > 0$ does not depend on n. Let W be the $V \times K$ matrix with for all (v, k), $W_{vk} = 1/\sqrt{|V_k|}$ if $v \in V_k$ and 0 otherwise. Let $W_\perp$ be an orthonormal basis of the space perpendicular to the linear span of W. The output $(S_k)_{1 \leq k \leq K}$ of the Detection algorithm with input (Q, V, K) satisfies: if $\|W_\perp^T Q\| = o(1)$, then there exists a permutation ζ of $\{1, \ldots, K\}$ such that*

$$\frac{\left|\bigcup_{k=1}^K S_k \setminus V_{\zeta(k)}\right|}{n} = O\left(\|W_\perp^T Q\|^2\right).$$

Step 1. We use the notations introduced in the pseudo-codes of the various algorithms. Let $M^{(g)} = \mathbb{E}[A^{(g)}]$ and $M' = \mathbb{E}[A']$. Let $A_\Gamma^{(g)} = (A_{vw}^{(g)})_{v,w \in \Gamma^{(g)}}$ and $M_\Gamma^{(g)} = (M_{vw}^{(g)})_{v,w \in \Gamma^{(g)}}$. Analogously, we define $A'_\Gamma = (A'_{vw})_{v,w \in \Gamma'}$ and $M'_\Gamma = (M'_{vw})_{v,w \in \Gamma'}$.

We prove that if $\frac{\sigma'_K}{\sqrt{m\hat{p}'}} \cdot 1_{\{m\hat{p}' \geq 50\}} = \omega(1)$, then the spectral method applied to the matrix of indirect edges A' is asymptotically accurate. We omit the proof of the asymptotic accuracy of the spectral method applied to $A^{(g)}$ under the condition $\frac{\sigma_K^{(g)}}{\sqrt{m\hat{p}^{(g)}}} \cdot 1_{\{m\hat{p}^{(g)} \geq 50\}} = \omega(1)$ (since it can be conducted in the same way).

Recall that σ'_K denotes the K-th largest singular value of the trimmed matrix A'_Γ. Observe that by assumption, for n large enough, $m\hat{p}' \geq 50$. Hence applying the law of large numbers, we can conclude that the largest singular value ξ_1 of A' scales at most as $m\hat{p}'$ w.h.p.. Since $\sigma'_K \leq \sigma'_1 \leq \xi_1$

(where σ'_1 is the largest singular value of A'_Γ) and $\frac{\sigma'_K}{\sqrt{m\hat{p}'}} = \omega(1)$, we deduce that $m\hat{p}' = \omega(1)$ w.h.p.. Hence the trimming step in the Approx algorithm applied to $(A', \hat{p}', V^{(g)}, K)$ does remove a negligible proportion of green nodes, i.e., w.h.p. $|V^{(g)} \setminus \Gamma| = o(|V^{(g)}|)$ or equivalently $|\Gamma'| = m(1 + o(1))$.

Observe that w.h.p., $\hat{p}' = \frac{\sum_{u,v} M'_{uv}}{m^2}(1 + o(1)) = \Theta(\max_{uv}\{M'_{uv}\})(1 + o(1))$ by the law of large numbers and $\sum_{w\in\Gamma'} A'_{vw} = O(m\hat{p}')$ for all $v \in \Gamma'$. From random matrix theory [8], with probability $1 - O(1/m)$, $\|A'_\Gamma - M'_\Gamma\| = O(\sqrt{m\hat{p}'})$. Next we apply Lemma 10 to (A'_Γ, Γ', K) and deduce that $\sigma'_K = \Theta(\lambda_K(A'_\Gamma))$ w.h.p.. From $\lambda_K(M'_\Gamma) \geq \lambda_K(A'_\Gamma) - \|A'_\Gamma - M'_\Gamma\|$, and $\frac{\sigma'_K}{\sqrt{m\hat{p}'}} = \omega(1)$, we deduce that w.h.p.,

$$\frac{\lambda_K(M'_\Gamma)}{\|A'_\Gamma - M'_\Gamma\|} = \omega(1).$$

If $M'_\Gamma = U\Lambda U^T$, we deduce from Lemma 11 applied to A'_Γ and M'_Γ that w.h.p., $\|U_\perp^T Q\| = o(1)$. We can now apply Lemma 12 replacing V by Γ' and V_k by $\Gamma' \cap V_k^{(g)}$. Observe that the linear span of U coincides with that of W (refer to Lemma 12 for the definition of W). Hence, w.h.p., the nodes Γ' are accurately classified, and so are the nodes in $V^{(g)}$.

Step 2. We distinguish two cases: 1. $\gamma f(n) = \omega(1)$; 2. $\gamma f(n) = O(1)$ and $\sqrt{\gamma} f(n) = \omega(1)$.

Case 1. Assume that $\gamma f(n) = \omega(1)$. By the law of large numbers, $m\hat{p}^{(g)} = \Theta(\gamma f(n))$ w.h.p.. Since $\lambda_K(M_\Gamma^{(g)}) = \Omega(\gamma f(n))$, $\|A_\Gamma^{(g)} - M_\Gamma^{(g)}\| = \Theta(\sqrt{m\hat{p}^{(g)}}) = \Theta(\sqrt{\gamma f(n)})$ and $\lambda_K(A_\Gamma^{(g)}) \geq \lambda_K(M_\Gamma^{(g)}) - \|A_\Gamma^{(g)} - M_\Gamma^{(g)}\|$, we get $\frac{\lambda_K(A_\Gamma^{(g)})}{\sqrt{m\hat{p}^{(g)}}} = \omega(1)$ w.h.p.. Since $\sigma_K^{(g)} = \Theta(\lambda_K(A_\Gamma^{(g)}))$ from Lemma 10, w.h.p.

$$\frac{\sigma_K^{(g)}}{\sqrt{m\hat{p}^{(g)}}} \cdot 1_{\{m\hat{p}^{(g)} \geq 50\}} = \omega(1).$$

Case 2. Assume that $\gamma f(n) = O(1)$ and $\sqrt{\gamma} f(n) = \omega(1)$. We first compute M'_{ij} for any $i, j \in V^{(g)}$. For notational simplicity, $\alpha_k = \frac{|V_k^{(g)}|}{n}$ and $\beta_k = \frac{|V_k^{(r)}|}{n}$.
(i) Let i, j be two green nodes belonging to the same community, i.e., $i, j \in V_k^{(g)}$. Let $v \in V_k^{(r)}$. We have:

$$\mathbb{P}\left[A_{vi} = 1 = A_{vj}, \sum_{w\in V^{(g)}} A_{vw} = 2\right] = p^2(1-p)^{\alpha_k n - 2} \prod_{l\neq k}(1-q)^{\alpha_l n}.$$

This probability is equivalent to $p^2 \exp(-\alpha_k pn - \sum_{l\neq k} \alpha_l qn)$ when $n \to \infty$. Similarly, when $v \in V_{k'}^{(g)}$ for some $k' \neq k$, the probability $\mathbb{P}[A_{vi} = 1 = A_{vj}, \sum_{w\in V^{(g)}} A_{vw} = 2]$ is equivalent to $q^2 \exp(-\alpha_{k'} pn - \sum_{l\neq k'} \alpha_l qn)$ when $n \to \infty$. We deduce that:

$$M'_{ij} \sim p^2 n\beta_k\eta_k + q^2 n \sum_{k'\neq k} \beta_{k'}\eta_{k'}, \quad \text{as } n \to \infty, \tag{8}$$

where $\eta_k = \exp(-\alpha_k pn - \sum_{l\neq k} \alpha_l qn)$.
(ii) Let i, j be two green nodes belonging to different communities, i.e., $i \in V_k^{(g)}$ and $j \in V_\ell^{(g)}$, for $k \neq \ell$. Using the same analysis as above, we have:

$$M'_{ij} \sim pq(\beta_k\eta_k + \beta_\ell\eta_\ell)n + q^2 n \sum_{k'\notin\{k,\ell\}} \beta_{k'}\eta_{k'} \quad \text{as } n \to \infty. \tag{9}$$

From (8)-(9) and the law of large numbers, we get w.h.p., $m\hat{p}' = \Theta(\gamma f(n)^2)$ (this comes from the facts that $\gamma f(n) = O(1)$ and $\alpha_k pn = \Theta(\gamma pn) = \Theta(\gamma f(n))$). As a consequence, $m\hat{p}' = \omega(1)$ w.h.p.. Thus, in the trimming process in the Approx algorithm applied to $(A', \hat{p}', V^{(g)}, K)$, we must have $|V^{(g)} \setminus \Gamma'| = o(|V^{(g)}|)$ w.h.p..

We also deduce from the above analysis that we can represent M'_Γ as follows:

$$M'_\Gamma = M^{(g)}_{\Gamma',K}\Lambda'(M^{(g)}_{\Gamma',K})^T,$$

where $M^{(g)}_{\Gamma',K}$ is a $\Gamma' \times K$ matrix where the k-th column of $M^{(g)}_{\Gamma',K}$ is the column vector of $M^{(g)}_\Gamma$ corresponding to $v \in V^{(g)}_k$, and Λ' is a $K \times K$ diagonal matrix where k-th element is $\beta_k\eta_k n$. Since $\frac{\|M^{(g)}_{\Gamma',K}\boldsymbol{x}\|}{\|\boldsymbol{x}\|} = \Omega(\sqrt{m}p)$ for any $\boldsymbol{x} \in \mathbb{R}^{K\times 1}$, $\lambda_K(M'_\Gamma) = \Omega(mnp^2 \min_{1\le k\le K}\eta_k) = \Omega(\gamma f(n)^2)$. By the law of large numbers, w.h.p., $m\hat{p}' = \Theta(\gamma f(n)^2)$. Then, as in the analysis of Case 1, we conclude that w.h.p.

$$\frac{\sigma'_K}{\sqrt{m\hat{p}'}} \cdot 1_{\{m\hat{p}'\ge 50\}} = \omega(1).$$

B.6 Proof of Lemma 10

To conclude $\sigma_K = \Theta(\lambda_K(A))$, we show that $\sigma_K = O(\lambda_K(A))$ (Step 1) and $\sigma_K = \Omega(\lambda_K(A))$ (Step 2).

Step 1. When $\|A\| = \Theta(\lambda_K(A))$, this is trivial since singular values of $R_{\tau^\star}$ have to be less than $\|A\|$. Let $\|A\| = \omega(\lambda_K(A))$. Then, there exists $\ell < K$ such that $\lambda_\ell(A) = \omega(\lambda_K(A))$ and $\lambda_{\ell+1}(A) = \Theta(\lambda_K(A))$. We denote by $\tilde{U}_j\tilde{\Lambda}\tilde{U}_j^T$ be the SVD of rank j approximation of A. Let $Q_{K,\tau^\star}$ denote the K-th column vector of $Q_{\tau^\star}$. Analogously with Step 1 of the proof of Lemma 11, we can show that $\|\tilde{U}_j^T Q_{K,\tau^\star}\| = O(\frac{\lambda_K(A)}{\lambda_j(A)})$ for all $j \le \ell$. Therefore, $\sigma_k = O(\lambda_K(A))$.

Step 2. When $\lambda_n(A) = \Theta(\lambda_K(A))$, this is trivial since singular values of $R_{\tau^\star}$ have to be larger than $\lambda_n(A)$. Let $\lambda_n(A) = o(\lambda_K(A))$. Then, there exists $\ell \ge K$ such that $\lambda_{\ell+1}(A) = o(\lambda_K(A))$ and $\lambda_\ell(A) = \Theta(\lambda_K(A))$. Analogously with Step 1 of the proof of Lemma 11, we can show that $\|(\tilde{U}_\ell)_\perp^T Q_{K,\tau^\star}\| = O(\frac{\lambda_{\ell+1}(A)}{\lambda_\ell(A)}) = o(1)$, where $(\tilde{U}_\ell)_\perp$ is an orthonormal basis of the perpendicular to the linear span of $\tilde{U}_\ell$. Therefore, $\sigma_k = \Omega(\lambda_\ell(A)) = \Omega(\lambda_K(A))$.

B.7 Proof of Lemma 11

We denote by $\tilde{A} = \tilde{U}\tilde{\Lambda}\tilde{U}^T$ be the SVD of rank K approximation of A. Let $U_\perp$ and $\tilde{U}_\perp$ be orthonormal bases of the perpendicular spaces to the linear spans of U and $\tilde{U}$, respectively. Since

$$\begin{aligned}\|U_\perp^T Q_{\tau^\star}\| &= \|U_\perp^T(\tilde{U}\tilde{U}^T + \tilde{U}_\perp\tilde{U}_\perp^T)Q_{\tau^\star}\| \le \|U_\perp^T\tilde{U}\tilde{U}^T Q_{\tau^\star}\| + \|U_\perp^T\tilde{U}_\perp\tilde{U}_\perp^T Q_{\tau^\star}\| \\ &\le \|U_\perp^T\tilde{U}\|\|\tilde{U}^T Q_{\tau^\star}\| + \|U_\perp^T\tilde{U}_\perp\|\|\tilde{U}_\perp^T Q_{\tau^\star}\| \le \|U_\perp^T\tilde{U}\| + \|\tilde{U}_\perp^T Q_{\tau^\star}\|,\end{aligned}$$

to conclude this proof, we will show that $\|\tilde{U}_\perp^T Q_{\tau^\star}\| = O\left(\frac{\|A-M\|}{\lambda_K(M)}\right)$ and $\|U_\perp^T\tilde{U}\| = O\left(\frac{\|A-M\|}{\lambda_K(M)}\right)$.

Step 1. $\|\tilde{U}_\perp^T Q_{\tau^\star}\| = O\left(\frac{\|A-M\|}{\lambda_K(M)}\right)$**:** Let $\boldsymbol{x}_1$ be the right singular vector of $\tilde{U}_\perp^T Q_{\tau+1}$ corresponding to the largest singular value and $\tilde{\boldsymbol{x}}_1$ be a $K \times 1$ vector such that $\boldsymbol{x}_1 = R_{\tau+1}\tilde{\boldsymbol{x}}_1$. Then,

$$\begin{aligned}\|\tilde{U}_\perp^T Q_{\tau+1}\|_2^2 &= \frac{\|\tilde{U}_\perp^T Q_{\tau+1}\boldsymbol{x}_1\|_2^2}{\|\boldsymbol{x}_1\|_2^2} = \frac{\|\tilde{U}_\perp^T Q_{\tau+1}R_{\tau+1}\tilde{\boldsymbol{x}}_1\|_2^2}{\|R_{\tau+1}\tilde{\boldsymbol{x}}_1\|_2^2} \\ &= \frac{\|\tilde{U}_\perp^T Q_{\tau+1}R_{\tau+1}\tilde{\boldsymbol{x}}_1\|_2^2}{\|\tilde{U}^T Q_{\tau+1}R_{\tau+1}\tilde{\boldsymbol{x}}_1\|_2^2 + \|\tilde{U}_\perp^T Q_{\tau+1}R_{\tau+1}\tilde{\boldsymbol{x}}_1\|_2^2} \\ &= \frac{\|\tilde{U}_\perp^T A Q_\tau\tilde{\boldsymbol{x}}_1\|_2^2}{\|\tilde{U}^T A Q_\tau\tilde{\boldsymbol{x}}_1\|_2^2 + \|\tilde{U}_\perp^T A Q_\tau\tilde{\boldsymbol{x}}_1\|_2^2} \\ &\le \frac{\|A-M\|_2^2}{(\lambda_K(M) - \|A-M\|_2)^2(1 - \|\tilde{U}_\perp^T Q_\tau\|_2^2) + \|A-M\|_2^2},\end{aligned} \tag{10}$$

where the last inequality stems from that

$$\begin{aligned}
\|\tilde{U}_\perp^T A Q_\tau \tilde{\boldsymbol{x}}_1\|_2 &\le \|\tilde{U}_\perp^T A\|\|Q_\tau\|\|\tilde{\boldsymbol{x}}_1\|_2 \le \|\tilde{U}_\perp^T A\|\|\tilde{\boldsymbol{x}}_1\|_2 = \lambda_{K+1}(A)\|\tilde{\boldsymbol{x}}_1\|_2 \le \|(A-M)\|\|\tilde{\boldsymbol{x}}_1\|_2 \\
\|\tilde{U}^T A Q_\tau \tilde{\boldsymbol{x}}_1\|_2 &= \|(\tilde{U}^T A\tilde{U}\tilde{U}^T Q_\tau + \tilde{U}^T A\tilde{U}_\perp\tilde{U}_\perp^T Q_\tau)\tilde{\boldsymbol{x}}_1\|_2 = \|\tilde{U}^T A\tilde{U}\tilde{U}^T Q_\tau \tilde{\boldsymbol{x}}_1\|_2 \\
&\ge \lambda_K(A)\|\tilde{U}^T Q_\tau \tilde{\boldsymbol{x}}_1\|_2 \ge (\lambda_K(M) - \|A-M\|)\|\tilde{U}^T Q_\tau \tilde{\boldsymbol{x}}_1\| \\
&\ge (\lambda_K(M) - \|A-M\|)\sqrt{1-\|\tilde{U}_\perp^T Q_\tau\|_2^2}\|\tilde{\boldsymbol{x}}_1\|.
\end{aligned}$$

Let $\zeta = \frac{\|A-M\|_2^2}{(\lambda_K(M)-\|A-M\|_2)^2}$. Since $\frac{\|A-M\|_2}{\lambda_K(M)} = o(1)$, $\zeta = O(\frac{\|A-M\|_2^2}{\lambda_K(M)^2}) = o(1)$. Then, from (10),

$$1 - \|\tilde{U}_\perp^T Q_{\tau+1}\|_2^2 \ge 1 - \frac{\zeta}{1-\|\tilde{U}_\perp^T Q_\tau\|_2^2+\zeta} = \frac{1-\|\tilde{U}_\perp^T Q_\tau\|_2^2}{1-\|\tilde{U}_\perp^T Q_\tau\|_2^2+\zeta}.$$

When $1-\|\tilde{U}_\perp^T Q_\tau\|_2^2 \le \zeta$, $1-\|\tilde{U}_\perp^T Q_{\tau+1}\|_2^2 \ge \frac{1-\|\tilde{U}_\perp^T Q_\tau\|_2^2}{2\zeta}$. From this, one can easily check that when $\tau \ge \frac{\log(\zeta/(1-\|\tilde{U}_\perp^T Q_0\|_2^2))}{\log(1/2\zeta)}$, $1-\|\tilde{U}_\perp^T Q_\tau\|_2^2 \ge \zeta$, $1-\|\tilde{U}_\perp^T Q_{\tau+1}\|_2^2 \ge 1/2$, and $1-\|\tilde{U}_\perp^T Q_{\tau+2}\|_2^2 \ge 1-2\zeta$. Therefore, when $\frac{\log|V|}{2} \ge \frac{\log(\zeta/(1-\|\tilde{U}_\perp^T Q_0\|_2^2))}{\log(1/2\zeta)}$, $\|\tilde{U}_\perp^T Q_{\tau^\star}\|_2 = O\left(\frac{\|A-M\|}{\lambda_K(M)}\right)$. Since $\zeta = o(1)$, to complete this proof, it is sufficient to show that $1-\|\tilde{U}_\perp^T Q_0\|_2^2 \ge 1/\text{Poly}(n)$ with probability $1-O(1/n)$, where Poly(n) is a polynomial function of n with finite order. By Theorem 1.2 of [16] (Please refer to the proof of Lemma 10 of [14]) we can conclude this part with Poly$(n) = 1/n^4$.

Step 2. $\|U_\perp^T\tilde{U}\| = O\left(\frac{\|A-M\|}{\lambda_K(M)}\right)$**:** We can get an upper bound and a lower bound for $\|AU_\perp\|$ as follows:

$$\begin{aligned}
\|AU_\perp\| &= \|(M+A-M)U_\perp\| = \|(A-M)U_\perp\| \le \|A-M\| \\
\|AU_\perp\| &= \|(\tilde{A}+A-\tilde{A})U_\perp\| \ge \|\tilde{A}U_\perp\| - \|(A-\tilde{A})U_\perp\| \ge \|\tilde{U}\tilde{\Lambda}\tilde{U}^T U_\perp\| - \|A-\tilde{A}\| \\
&\ge \lambda_K(A)\|\tilde{U}^T U_\perp\| - \|A-\tilde{A}\| \ge (\lambda_K(M) - \|A-M\|)\|\tilde{U}^T U_\perp\| - \|A-M\|.
\end{aligned}$$

When we combine above bounds, $\|\tilde{U}^T U_\perp\| \le \frac{2\|A-M\|}{(\lambda_K(M)-\|A-M\|)} = O\left(\frac{\|A-M\|}{\lambda_K(M)}\right)$.

B.8 Proof of Lemma 12

From the definitions of W and $W_\perp$, $Q = WW^TQ + W_\perp W_\perp^T Q$. Since rows of W corresponding to the nodes from the same cluster are the same, the rows of WW^TQ are also the same for the node from the same clusters. Let $WW^TQ(k)$ be the rows of WW^TQ corresponding to $v \in V_k$. Let $\boldsymbol{v}_{(k\ell)} \in \mathbb{R}^{K\times 1}$ such that k-th row and ℓ-th row are $1/\sqrt{|V_k|}$ and $-1/\sqrt{|V_\ell|}$, respectively and other elements are zero. Then, $\|WW^TQ(k) - WW^TQ(\ell)\|^2 = \|W^TQ\boldsymbol{v}_{(k\ell)}\|^2$. Since $\|W^TQ\boldsymbol{x}\| \ge \sqrt{1-\|W_\perp^TQ\|^2}\|\boldsymbol{x}\|$,

$$\|WW^TQ(k) - WW^TQ(\ell)\|^2 = \Omega(\frac{1-\|W_\perp^TQ\|^2}{n}) = \Omega(\frac{1}{n}) \quad \text{for all } k \ne \ell.$$

Therefore, with some positive $C > 0$,

$$\begin{aligned}
C\frac{|\bigcup_{k,\ell:k\ne\ell} S_k \bigcap V_\ell|}{n} &\le \sum_{k,\ell:k\ne\ell}\sum_{v\in S_k\bigcap V_\ell} \|WW^TQ(k) - WW^TQ(\ell)\|^2 \\
&\le 2\sum_{k,\ell:k\ne\ell}\sum_{v\in S_k\bigcap V_\ell} \|WW^TQ(k) - \xi_{i^\star,k}\|^2 + \|\xi_{i^\star,k} - WW^TQ(\ell)\|^2 \\
&\le 4\sum_{k,\ell:k\ne\ell}\sum_{v\in S_k\bigcap V_\ell} \|WW^TQ(\ell) - \xi_{i^\star,k}\|^2 \\
&\le 8\sum_{k,\ell:k\ne\ell}\sum_{v\in S_k\bigcap V_\ell} \|WW^TQ(\ell) - Q_v\|^2 + \|Q_v - \xi_{i^\star,k}\|^2 \\
&\le 8\|W_\perp W_\perp^TQ\|_F^2 + 8r_{i^\star} \le 8K\|W_\perp W_\perp^TQ\|^2 + 8r_{i^\star} \le 8K\|W_\perp^TQ\|^2 + 8r_{i^\star}
\end{aligned}$$

To conclude this proof, we need to show that $r_{i^\star} = O(\|W_\perp^T Q\|^2)$. Let i^t be an integer between 1 and $\log n$ such that $\frac{\|W_\perp W_\perp^T Q\|_F^2}{n\delta^2} \leq \frac{i^t}{n\log n} \leq \frac{\delta^2}{n}$ with positive constant δ close to 0. There exists such i^t for any δ, since $\|W_\perp W_\perp^T Q\|^2 = o(1)$ and the rank of $W_\perp W_\perp^T Q$ is K. Then,

$$\Big|\bigcup_{1\leq k\leq K}\{v\in V_k : \|Q_v - WW^TQ(k)\|^2 \leq \frac{i^t}{4n\log n}\}\Big| \geq n - \|W_\perp W_\perp^T Q\|_F^2 \frac{4n\log n}{i^t}$$
$$\geq n(1-4\delta^2).$$

From this, since $\|Q_v - Q_w\|^2 \leq 2\|Q_v - WW^TQ(k)\|^2 + 2\|Q_w - WW^TQ(k)\|^2$, when v satisfying that $\|Q_v - WW^TQ(k)\|^2 \leq \frac{i^t}{4n\log n}$,

$$|X_{i^t,v}| \geq |V_k| - 4\delta^2 n.$$

On the other hand, since $\|Q_v - Q_w\|^2 \geq \frac{1}{2}\|Q_v - WW^TQ(k)\|^2 - \|Q_w - WW^TQ(k)\|^2$, when v satisfying that $\|Q_v - WW^TQ(k)\|^2 \geq \frac{i^t}{n\log n}$,

$$|X_{i,v}| \leq 4\delta^2 n.$$

With small enough constant δ, therefore, when v and w satisfy that $\|Q_v - WW^TQ(k)\|^2 \leq \frac{i^t}{4n\log n}$ and $\|Q_w - WW^TQ(k)\|^2 \geq \frac{i^t}{n\log n}$, $|X_{i,v}| > |X_{i,w}|$, which indicates that the origin of $T_{i^t,k}$ is at least $\|Q_{v_k} - WW^TQ(k)\|^2 \leq \frac{i^t}{n\log n}$ and $|T_{i^t,,k}| \geq |V_k| - 4\delta^2 n$. Since $\|\cdot\|$ is a convex function, by Jensen's inequality, for all k,

$$\|WW^TQ(k) - \xi_{i^t,k}\|^2 \leq \frac{\sum_{v\in T_{i^t,k}}\|WW^TQ(k) - Q_v\|^2}{|T_{i^t,k}|} \leq \frac{\|W_\perp W_\perp^T Q\|_F^2}{|V_k| - 4\delta^2 n} = O(\frac{\|W_\perp^T Q\|^2}{n}).$$

Therefore,

$$r_{i^t} = \sum_{k=1}^K \sum_{v\in T_{i^t,k}} \|Q_v - \xi_{i^t,k}\|^2 \leq \sum_{k=1}^K \sum_{v\in V_k} \|Q_v - \xi_{i^t,k}\|^2$$
$$\leq 2\sum_{k=1}^K \sum_{v\in V_k} \|Q_v - WW^TQ(k)\|^2 + \|WW^TQ(k) - \xi_{i^t,k}\|^2$$
$$\leq 2\|W_\perp W_\perp^T Q\|_F^2 + 2\sum_{k=1}^K \sum_{v\in V_k} \|WW^TQ(k) - \xi_{i^t,k}\|^2$$
$$= O(\|W_\perp^T Q\|^2) + 2\sum_{k=1}^K \sum_{v\in V_k} \|WW^TQ(k) - \xi_{i^t,k}\|^2 = O(\|W_\perp^T Q\|^2).$$

Since $r_{i^\star} \leq r_{i^t}$, $r_{i^\star} = O(\|W_\perp^T Q\|^2)$.

B.9 Proof of Theorem 5

Let $\mu(v, S_k^{(g)}) = \mathbb{E}[\sum_{w\in S_k^{(g)}} A_{vw}]$ and $Var(v, S_k^{(g)}) = \mathbb{E}[(\mu(v, S_k^{(g)}) - \sum_{w\in S_k^{(g)}} A_{vw})^2]$. Since $|\bigcup_{k=1}^K (S_k^{(g)} \setminus V_k^{(g)})| = o(|V^{(g)}|)$ from Theorem 4, $\frac{\mu(v,S_k^{(g)})}{|S_k^{(g)}|} = p(1+o(1))$ and $\frac{Var(v,S_k^{(g)})}{|S_k^{(g)}|} = p(1+o(1))$ when $v\in V_k$, and $\frac{\mu(v,S_k^{(g)})}{|S_k^{(g)}|} = q(1+o(1))$ and $\frac{Var(v,S_k^{(g)})}{|S_k^{(g)}|} = q(1+o(1))$ when $v\notin V_k$. By Chebyshev's inequality, when $v\in V_k$, $v\in S_k$ with high probability since $\frac{\mu(v,S_k^{(g)}) - \mu(v,S_{k'}^{(g)})}{\sqrt{Var(v,S_k^{(g)})}} = \omega(1)$ for all $k'\neq k$ when $\gamma f(n) = \omega(1)$.

B.10 Proof of Theorem 6

In this proof, we use Chernoff bound as the form of Lemma 8.1 in [5].

From Theorem 4, Algorithm 1 classifies the arrival nodes at each time block with diminishing fraction of misclassified nodes. Between $S_i^{(\tau)}$ and $S_j^{(\tau+1)}$, the number of connections is $\Theta(B^2\frac{\bar{f}(n)}{n}) = \Theta(\frac{h^2(n)n}{\min\{f(n),n^{1/3}\}\log^2 n}) = \omega(1)$ from the condition of this theorem. Let $\mu(k,i) = \frac{\sum_{v\in\hat{V}_i}\sum_{w\in S_k^{(\tau)}} A_{wv}}{\sum_{v\in\hat{V}_i}\sum_{w\in S_k^{(\tau)}} 1}$. By the Chernoff bound, with high probability (since $\sum_{v\in\hat{V}_i}\sum_{w\in S_k^{(\tau)}} A_{wv} = \omega(1)$), $\mu(k,i) = p(1-o(1))$ when $\frac{|S_k^{(\tau)}\bigcap V_i|}{|S_k^{(\tau)}|} = 1-o(1)$ and $\mu(k,i) = q(1+o(1))$ when $\frac{|S_k^{(\tau)}\bigcap V_i|}{|S_k^{(\tau)}|} = o(1)$. Therefore, with high probability, $S_k^{(\tau)}$ is merged with $\hat{V}_{s(k)}$ such that $\frac{|S_k^{(\tau)}\bigcap V_{s(k)}|}{|S_k^{(\tau)}|} = 1-o(1)$. Thus, $\frac{|\hat{V}_k\bigcap V_k|}{|\hat{V}_k|} = 1-o(1)$ for all k with high probability.

Since $\frac{|\hat{V}_k\bigcap V_k|}{|\hat{V}_k|} = 1-o(1)$ for all k, one can easily show using the Chernoff bound that $\frac{N_{v,k}}{|\hat{V}_k|} \geq p(1-\frac{p-q}{4})$ when $v\in V_k$ and $\frac{N_{v,k'}}{|\hat{V}_{k'}|} \leq q(1+\frac{p-q}{4})$ when $v\notin V_{k'}$ with probability $1-O(\exp(-cT\frac{f(n)}{n}))$ with a constant $c>0$. Thus, the probability for that $\frac{N_{v,k}}{|\hat{V}_k|} \leq \frac{N_{v,k'}}{|\hat{V}_{k'}|}$ for $v\in V_k$ and $k\neq k'$ is $O(\exp(-cT\frac{f(n)}{n}))$.